Enhancing the Scientific Process with Artificial Intelligence: Forest Science Applications

Ronald E. McRoberts[1]
Daniel L. Schmoldt[2]
H. Michael Rauscher[3]

Abstract. *Forestry, as a science, is a process for investigating nature. It consists of repeatedly cycling through a number of steps, including identifying knowledge gaps, creating knowledge to fill them, and organizing, evaluating, and delivering this knowledge. Much of this effort is directed toward creating abstract models of natural phenomena. The cognitive techniques of AI, with their emphasis on knowledge and thinking, can help scientists create, manipulate, and evaluate these models. The steps of the scientific process can be enhanced with five cognitive techniques from AI: neural networks, machine learning, advisory systems, knowledge management, and qualitative simulation. For each technique, we identify the steps of the scientific process to which it can be applied, provide background for the technique, and identify current or potential applications in forestry.*

Scientific research has been construed as both a problem-solving activity (Kleiner 1985) and as a process for producing knowledge. One component of the scientific process is identifying problems and knowledge gaps. However, the distinguishing features of science relate to the methods it uses to solve the problems and to create the knowledge to fill the gaps. Mario Bunge (1967) asserts that science is distinguished by its unique goal and its unique method. He claims that the goal of science is to map patterns of facts and that the scientific method is "a mark of science . . . no scientific method, no science." At the very least, the scientific method is a means of solving problems and creating knowledge that involves proposing solutions or hypotheses and then testing them for adequacy.

However, the scientific process consists of more than just identifying and solving problems, creating knowledge, and collecting the results.

[1]USDA Forest Service, North Central Forest Experiment Station, St. Paul, Minnesota 55108.
[2]USDA Forest Service, Southeastern Forest Experiment Station, Thomas M. Brooks Forest Products Center, Virginia Polytechnic Institute and State University, Blacksburg, Virginia 24061.
[3]USDA Forest Service, North Central Forest Experiment Station, Grand Rapids, Minnesota 55744.

The solutions and created knowledge must be evaluated for consistency and completeness and systematically and coherently organized. A usual by-product of the evaluations and organization is identification of additional problems and knowledge gaps. There is an additional component in the scientific process; the knowledge that has been created, evaluated, and organized must also be delivered to end users. While there is controversy as to whether the delivery of knowledge is a legitimate activity for researchers, it is nevertheless a necessary activity. Thus, we perceive science to be a cyclical process that includes identifying problems and knowledge gaps, proposing and testing solutions and hypotheses, evaluating the results, organizing the survivors, identifying new problems and knowledge gaps, and delivering the results (Fig. 1). Our intent is to discuss this process in more detail and to discuss how the cognitive aspects of artificial intelligence might or can be applied to enhance the process in forest science.

The Scientific Process

Knowledge gaps. A problem might be defined as a situation in which there are two states of knowledge, the present state and the desired state, with the difference between them representing a knowledge gap. Knowledge gaps are constructs that depend on the perspective of the researcher and currently accepted paradigms, and that exist in the form of interconnected sets in such a way that change in the state of one changes the states of the others. One way to identify knowledge gaps seems to be to attempt to model the desired state of knowledge. Inevitably, knowledge gaps become apparent as by-products of the modeling process. Nevertheless, there is very little information about how to define knowledge gaps in an efficient and effective manner.

Knowledge creation. The scientific method is the primary means of creating knowledge in a scientific discipline. In discussing the components of scientific method, Hans Reichenbach (1938) emphasized the distinction between the context of discovery and the context of justification. Discovery begins with a problem or knowledge gap, and ends with one or more proposed hypotheses, models, or solutions. Justification begins with the set of hypotheses, models, or solutions, and ends with justifiable inferences concerning them. Discovery has traditionally been considered a creative enterprise for which no logic could be constructed, while justification has a fairly well-defined inferential logic.

Numerous strategies have been used to discover hypotheses, including:

Trial and Error
Systematic Search
Serendipity
Inspiration
Analogy
Derivation
Illumination through Preparation
Induction.

A recent innovation that will be discussed in more detail later, and which may ultimately lead to a logic of discovery, involves applying the techniques of AI to discovering hypotheses. The consensus at this point, however, is that there are no right or wrong ways to achieve discovery; in essence, anything goes!

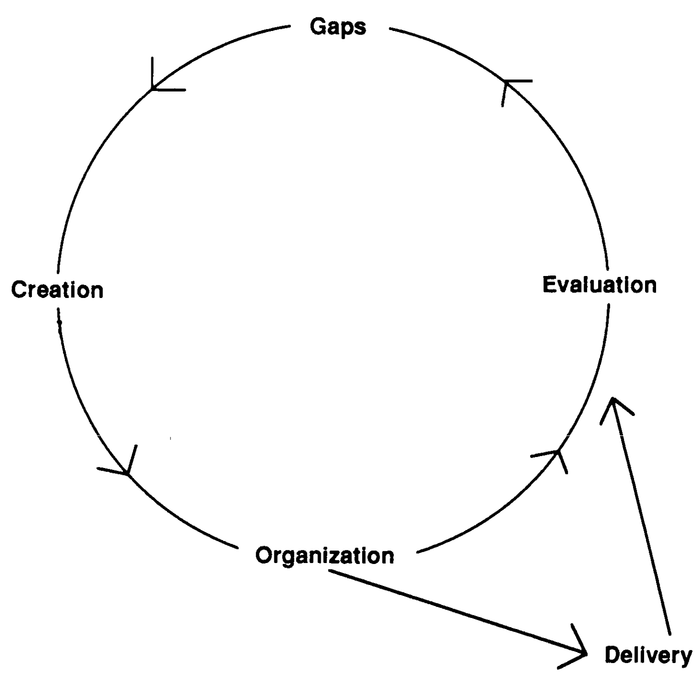

Figure 1. *The cyclical process of science.*

However discovered, hypotheses must be testable in the sense of being sensitive to evidence. If false, it must be possible to obtain contradicting empirical evidence. If true, it must be possible to obtain corroborating empirical evidence. For example, it would be virtually impossible to empirically corroborate a universal hypothesis, such as all swans are white, and likewise it would be virtually impossible to empirically contradict an existence hypothesis, such as there exist signals that travel faster than the speed of light.

While anything goes in discovery, it is quite different in justification. There exists a logic of justification and it is much more rigorous than is the methodology of discovery. It has not always been so. It was only in the 17th century that rigor began to be introduced into science. In his fundamental treatise on the logic of scientific method, Francis Bacon (1620) insisted upon a gradual passage from concrete facts to broad generalizations, and upon the use of controlled experimentation, not just observation. The phenomenal successes of Newton, who used and extended Bacon's methods, firmly established empiricism as a fundamental principle of science. However, in his exaltation of induction and experiment, Bacon also held that general laws could be established with complete certainty by using these almost mechanical processes. It was not until the 18th century that David Hume (1739) debunked the myth of scientific proof by inductive methods. With empirical and experimental methods in hand and a clear understanding of the impossibility of proof by inductive methods, modern science emerged.

From the time of the emergence of modem science until the mid-20th century, corroborating hypotheses was the primary strategy for creating new knowledge. Although scientists acknowledged corroboration was not proof, it was generally conceded to be the best that could be attained. One of the more powerful variations of this strategy is called hypothetico-deduction. With this strategy, a hypothesis is discovered by any means available to the researcher. From the hypothesis, a prediction or deduction is derived which is then compared to empirical evidence. If the comparison is favorable, the evidence is regarded as corroborating the hypothesis.

In the 1930s, Karl Popper (1935) suggested that the difficulties due to the failure of corroboration strategies to provide proof could be avoided by shifting to falsification strategies. Popper argued that conclusive disproof is possible because it takes only a single counterexample to disprove a hypothesis. Thus, falsification strategies attempt to gather the specific kind of evidence that contradicts hypotheses. Popper contends that science advances by disproof because hypotheses are conclusively eliminated from further consideration. The only results regarded as corroborating evidence for hypotheses are new and interesting failures to detect counterexamples where they would be most expected to occur.

All variations of corroboration and falsification strategies possess some degree of rigor. For each strategy, the researchers are obligated to convince their peers that the data acquisition methods are legitimate, objective, and untainted, and that the inference regarding the corroboration or falsification of the hypothesis is valid. While the method is irrelevant in discovery, in justification, the method is critical, must be open to scrutiny, and has considerable bearing on whether the inference will be accepted. For a more complete discussion of research strategies, see McRoberts (1989).

Knowledge organization. A collection of corroborated hypotheses is not particularly useful until it has been systematized into a body of knowledge. Because scientists use the terms hypothesis, law, and theory ambiguously, we will establish working definitions for this discussion. A hypothesis is a statement that refers to a pattern of facts and whose truth is as yet undetermined; a law is a well-corroborated hypothesis concerning the pattern of an entire class of facts; and a theory is a system of related hypotheses including some at the law level. Bunge (1967) includes the following as purposes for constructing theories: 1) to offer a map or model of a segment of reality, 2) to systematize knowledge by establishing logical relations among previously disconnected facts and hypotheses, and 3) to explain facts by means of systems of hypotheses from which they may be logically deduced. Thus, the search for theories is the search for systematic frameworks for bodies of knowledge.

Knowledge evaluation. Knowledge must be carefully evaluated, not only for empirical corroboration, but also for conceptual and logical consistency with existing knowledge. First, the concepts in hypotheses that are candidates for inclusion into a theory must be consistently defined with respect to those already included in the theory. Second, candidate hypotheses must not contradict hypotheses already in the theory, and theories augmented by new hypotheses must not generate mutually contradictory deductions. Third, new hypotheses must be evaluated for their effects on bridging or creating additional knowledge gaps within theories. Fourth, new hypotheses must be evaluated to determine if including them in a theory makes it possible to deduce or explain previously anomalous facts.

Knowledge delivery. Historically, scientists have not agreed as to whether the scientific process should include the delivery of knowledge to users. Traditionally, most scientists have been content to deliver their findings to other scientists via scholarly journals. The responsibility for ferreting out new knowledge and applying it has typically fallen upon intermediary agents or upon users.

Two recent trends have been influential in beginning to alter this tradition. First is the trend among some scientific institutions, including USDA Forest Service Research, to reward scientists for transferring knowledge to users. Second is the advent of knowledge-based advisory systems. The latter development has opened an entirely new means of delivering knowledge, has been widely and enthusiastically embraced by both creators and users of knowledge, and seems to hold great promise for effective and efficient knowledge delivery.

In summary, we view science as a cyclical process that includes a) identifying knowledge gaps, b) applying the scientific method to create new knowledge,c) systematically organizing knowledge through the search for laws and theories, d) evaluating newly created knowledge for inclusion into existing theories, and e) delivering knowledge to those who apply it. With this view in mind, we will now discuss some of the methods of AI to show how they are or might be applied to enhance the scientific process.

Cognitive Aspects of Artificial Intelligence

AI methods have received much attention at all levels in the computer science community. One of the first questions usually asked is: “What is AI?” In keeping with the scientific focus of this manuscript, we claim that AI is the study of knowledge production. This includes defining knowledge, describing its creation, and understanding how it is used. Knowledge may refer to academic and technical subjects or to something as apparently mundane as how one understands a joke. Knowledge production refers to a process analogous to what might occur in manufacturing beginning with an order for some product. The necessary inventories of component parts are collected or produced, assembled into a final product, and then, before delivery, the product is inspected for its quality. As we better understand how to create, manipulate, organize, and transmit knowledge, we become more efficient producers, distributors, and consumers of knowledge goods.

For purposes of illustrating AI applications to forestry science, we will focus our discussion on AI topics that deal with knowledge and thinking. Behavioral adaption, sensory interpretation, and physical interaction topics are also active areas of AI research. However, they seem less immediately important to forestry science than are knowledge and thinking, so we will not pursue those topics here. Because science is largely a thought process, we will limit our discussion to those aspects of AI that address issues of cognition. Cognitive activities are those that deal with modeling and understanding mental processes, with little reference to interactions with the outside world. In very general terms, this amounts to the creation, manipulation, organization, and translation of concepts and ideas.

Ideas exist as collections of facts and information drawn from the world around us (Hofstadter 1979, Johnson 1986). Synthesizing ideas from information produces knowledge (Michie and Johnston 1985, Rauscher and Schmoldt 1990). Synthesis may involve recognizing patterns, forming concepts, structuring and manipulating information and ideas, and converting and transmitting ideas and information to other agents. Epistemologically, if the process of science is the production of knowledge, the analysis of its assembly line remains one of the primary goals of AI.

* * *

We will examine the following AI methods: neural networks, machine learning, advisory systems, knowledge management, and simulation. The first two topics address the discovery of patterns. Advisory systems and knowledge management systems help structure and organize what is known and also help elucidate knowledge gaps. Heuristic simulation/quantitative modeling of knowledge aids the explanation and organization of current knowledge, thereby creating a stronger paradigm to support further scientific investigation. Research efforts in these areas can have substantial impact on understanding human thought processing in general, and also on how we think and what we know in the specialty that is forest science.

A number of notions regarding knowledge processing appear over and over in the literature and are regarded as fundamentals of AI at this early stage of its development. These are: knowledge representation, reasoning, search, and pattern recognition. In the remainder of this section, we examine those recurring ideas as preparation for a more detailed discussion of the previously identified AI methods.

Knowledge representation. Knowledge representation continues to be a critical component of any discussion about knowledge processing. We beg the question of what exactly constitutes knowledge and refer the reader to Barr and Feigenbaum (1981), Dretske (1986), and Giere (1984) for a variety of interpretations. However, knowledge is often viewed as an entity more refined than either data or information (Rauscher and Schmoldt 1990, Tanimoto 1987). Stillings et al. (1987) define a knowledge representation scheme as a symbol system composed of formal conventions (syntax) and some method for interpreting those conventions (semantics). Terms analogous to those used in linguistics are appropriate because we are dealing with a knowledge "language" used to express and communicate thoughts and ideas. If we want to make anthropomorphic comparisons, then knowledge representation shall include the capacities to represent new knowledge in addition to old, to "memorize" knowledge when encountered, to "recall" previously memorized knowledge, and to "think" about knowledge. Luger and Stubblefield (1989) describe knowledge as active, rather than passive, containing specifications for its use. Because there are so many different ways to represent knowledge (Tanimoto 1987), it is implied that people possess an extensive ability to represent what they know in diverse ways. All of the cognitive areas of AI mentioned above address some aspect of acquiring, retrieving, or reasoning about the knowledge inherent within a representation scheme.

Reasoning. Often we want to produce new knowledge, and to do so based upon what we already know or don't know. Without some capability to reason about what is or isn't known, knowledge functions only like a book, that is, a static account of some subject. It takes the reasoning abilities and the understanding of a reader to give it meaning and functionality. This is the reasoning component of knowledge representation mentioned above. Reasoning also includes control, or navigation, of a knowledge representation scheme. This meta-level of knowledge is often referred to as meta-knowledge.

Numerous models of reasoning are used in AI systems. Probably the most widely used is deductive reasoning. However, Stillings et al. (1987) note that there is substantial evidence to suggest that this is not how many people actually think, and, in fact, most people are very poor at it. On the surface it seems that deductive reasoning may bean important part of the scientific method, although there is much evidence for other forms of reasoning. Inductive reasoning was investigated extensively by Langley et al. (1987). Abductive reasoning (hypothesizing an antecedent based upon the truth of a conclusion) has been found quite common and useful in diagnostic problem-solving situations (Reggia 1985, Schmoldt and Martin 1989). A variety of reasoning methods have been developed in the general area of inexact reasoning, that is, reasoning with incomplete and uncertain knowledge. Some of these are: fuzzy reasoning (Zadeh 1965), belief networks based upon Bayes theorem (Pearl 1986), and the Dempster-Shafer theory of evidence (Shafer 1976). Formal schooling emphasizes the importance of logical reasoning, particularly in science. So it seems quite surprising that some of our most creative and innovative discoveries have occurred through thinking processes that are not based on classical logic (Halpern 1989). These alogical mechanisms often enhance our creativity by allowing us to juxtapose seemingly unrelated ideas (Halpern 1989).

Analogical reasoning is also a useful thinking method (Halpern 1989, Stillings et al. 1987). It is often possible to understand one phenomenon by comparing it to another phenomenon possessing some similar characteristics, then additional characteristics of the phenomenon of interest are hypothesized in a homomorphic fashion. Because much science involves the life and death of theories, we continually revise our beliefs in light of newly acquired knowledge. This nonmonotonic nature of scientific belief over time has also been an important consideration for AI systems: consequently, much work recently has focused on nonmonotonic reasoning (Doyle, 1979, McCarthy 1980, McDermott 1980, Reiter 1980).

Search. Given that an appropriate representation of knowledge exists for some subject, we often want to use it to solve problems. Problem solving can be viewed as finding a set of facts and hypotheses which describe a path to some state of affairs with desirable properties. This path can be found by searching through some subset of the pertinent knowledge. Search involves deciding what to do next, or where to look for a solution. For many small problems, exhaustive search works very well; it is guaranteed to find the best solution, if one is represented in the knowledge. Larger problem spaces often require more informed, heuristic searches to find a solution efficiently.

Heuristics are short cuts, good guesses, or tricks for evaluating alternatives within a problem representation. A heuristic may be any doma.in-specific knowledge that allows one to solve problems more efficiently in that domain. They generally arise through experience as a way of quickly and efficiently categorizing large quantities of information. The price paid for this simplicity is accuracy. By nature, heuristics are fallible and are not guaranteed to work correctly in all cases. Heuristics are often classified as strong or weak. Strong heuristics contain a large amount of subject area knowledge and hence are only applicable to a limited class of problems. Weak heuristics are applicable to a wider variety of problems. Because weak heuristics lack a narrow focus, they do not search as effectively as the stronger variety. A large amount of experimental evidence indicates that much of our problem solving is heuristic, rather than deductive or some other formalized method (Stillings et al. 1987). Langley et al. (1987) created a number of computer programs to perform scientific discovery using weak heuristics, so it is certainly not impossible that similar processes operate when implementing the scientific method.

Pattern recognition. Science seeks to describe and explain regularities within a changing and complex world. Recognizing these regular pattems helps with the creation and organization of knowledge. Hence, pattern recognition surfaces as an important part of the scientific process. Similarly, AI’s interest in pattern recognition stems from its importance for modifying behavior in response to past experience, i.e., its learning ability. Pattern recognition may involve creating new categories or classes of regularities that describe a set of observations, or it may attempt to relate conceptually two sets of patterns that are related physically, thereby modeling a physical relationship conceptually. Most AI systems learn very poorly. Some programs are designed specifically as learning programs; most others must be specifically told what knowledge to add and when to add it to their representation. Until AI systems can modify themselves, their “intelligence” will always seem more encyclopedic than adaptive, despite the very best knowledge representation schemes, search methods, and heuristics.

Neural Networks

Many of the models and techniques developed in AI borrow from naturally occurring models of complex system organization. The human brain is often used as a model, but cultural/social systems and biological systems (e.g., an ant colony) have provided valuable insights for AI. Neural networks, also referred to as parallel distributed processing (PDP) by Rumelhart and McClelland (1986), follow this natural/artificial association quite closely by modeling brain nerve cells and their interconnections. Despite some successes, production system models of intelligent behavior have been criticized on theoretical as well as practical grounds (Allman 1989). The production system paradigm represents a top-down approach to reasoning. This perspective assumes that reality can be described by abstract symbols that represent objects and their relationships to one another. Thought is based on symbols. Indeed, human consciousness is almost synonymous with language, the mental manipulation of symbols (Smith 1985).

While high-level, abstract thinking can often be reasonably formulated in this way (e.g., expert system knowledge), this symbolic model begins to break down when applied to sensory level activities such as speech recognition, vision, and pattern recognition. In response to the limitations of symbolic processing, some AI researchers have turned to recent advances in neurophysiology in their search for more useful models of some aspects of human thought processing. These models signify a bottom-up description of thinking. Neural network methodology and the traditional symbolic AI methods should be regarded not as rivals but as cooperators. The strengths of each compensate for the weaknesses of the other. Neural networks operating at the sensory level, integrated with symbolic AI at the cognitive level, promise to advance our ability to create artificial intelligence (Narasimhan 1990).

Network architectures. An artificial neuron is depicted in Figure 2a. In general, a neuron environment consists of an input vector $\mathbf{X} = \mathbf{X}_1, \mathbf{X}_2, \ldots . \mathbf{X}_n$, a weight vector **W,** a summation block that combines the inputs and weights, and an activation function f that transforms the weighted sum into an output signal. The input and weight vectors are usually numeric quantities, but need not be so (Rumelhart and McClelland 1986). The net result of applying the weights **W** to the inputs **X** is the value, NET= **XW.** An activation function **f,** usually a nonlinear function or a threshold function, is applied to the value NET to produce an output value, OUT= f(NET). Nonlinearity in the activation function more closely represents transfer characteristics of biological neuron activation (Wasserman 1989). Several of these neurons may be placed within a single layer of a network to create what is often called a perception (Fig. 2b). Minsky and Papert (1969) found, however, that perceptions were extremely limited in representing some fairly simple operations, e.g., exclusive-or. When perceptions are organized into multiple-layer networks (Fig. 2c) and utilize a nonlinear activation function, they overcome single-layer limitations and become quite powerful computationally (Wasserman 1989).

The last illustration (Fig. 2c) represents the general architecture of a multiple-layer neural network. Many alternative structures have also been developed for special purposes (Wasserman 1989). For a network to operate correctly, it must be trained by adjusting the weight arrays. There are two general types of training methods: supervised and unsupervised. In supervised training, each input vector is paired with a target output vector to produce a training pair. As each input vector is applied to the network, its output is compared to the target and an error is calculated. This error is used to adjust the weights so as to minimize the error. When the errors for the entire training set are acceptably low, then the network has been trained. Unsupervised training requires no target vectors, only input vectors. The training algorithm modifies weight vectors to produce output vectors that are consistent. Hence, similar input vectors are organized into classes designated by the output patterns. In addition, some

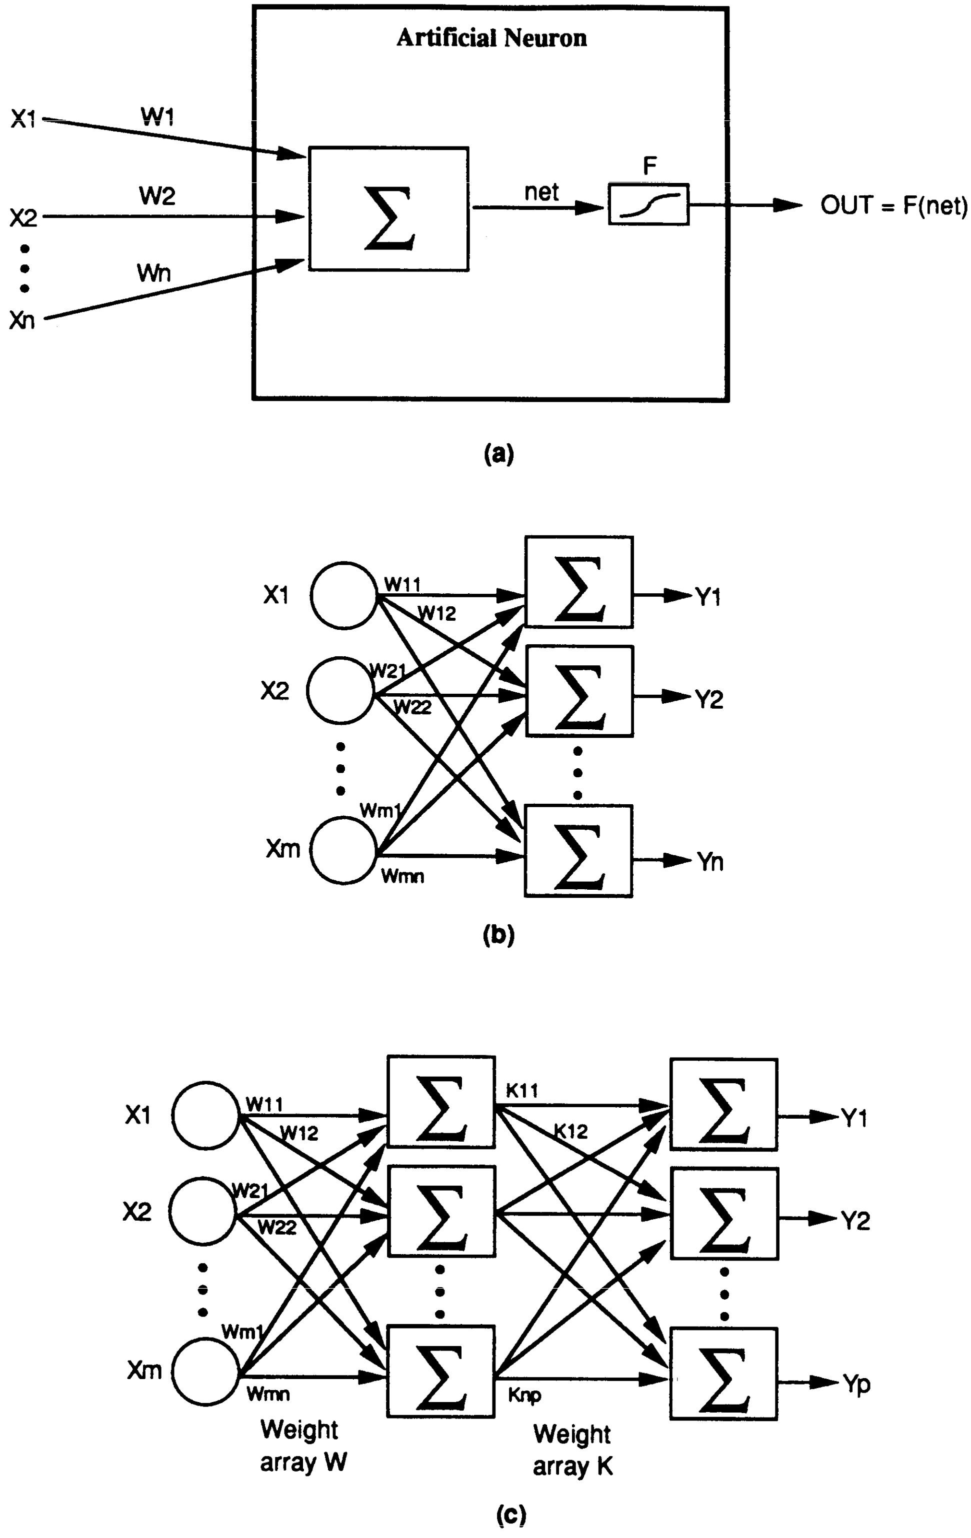

Figure 2. *A single artificial neuron (a) contains inputs, weights, summation function, and activation function to produce an output, OUT. A number of these processing units can be arranged into a single layer network (b), termed a perception. More complex and useful networks contain multiple layers (c).*

transformation may need to be applied to convert the output patterns into intuitively understandable classes. Just as there are many different architectures for neural nets, there are also numerous different training algorithms that can be applied to them.

This diversity in architectures and training methods permits a rich and varied repertoire of applications. Neural networks have been applied to pattern recognition in computer vision (Fukushima 1986), handwriting recognition (Burr 1987), data compression (Hecht-Nielsen 1988), analog-to-digital conversion (Tank and Hopfield 1986), and optimization (Hopfield and Tank 1985). As new network structures are created and new algorithms are applied to them, this list of applications will continue to grow.

Applications to science. Neural networks are adept at associating one pattern with another or correlating input with output. As such they are able to help with the creation of knowledge, even with noisy data. That is, an input pattern need not be identical to a previously learned pattern for it to be recognized and classified correctly. Classifying patterns is one way of discovering relationships between objects/events and their properties. This idea has been exploited previously by using statistical methods. In fact, White (1989) shows that the popular, backpropagation learning rule is analogous to nonlinear least squares regression. In addition to categorizing different objects/events, one can also associate one object/event with another by treating one as an input vector and the other as an output vector. It may be possible to attach reasonable, conceptual descriptions (in meaningful terms of the subject area) to the internal structure of a neural network. This may be accomplished by identifying the salient features of an output vector and tracing back through the network over large positive or negative weights, eventually returning to particular features of the input vector (Caudill 1989). Both pattern classification and pattern association use information about instances to generalize to a pattern; this is often referred to as induction or empirical discovery.

Certainly, neural networks, once developed, provide a mechanism for transferring research results to end users. Once trained for a particular problem set, these networks can help people solve problems. However, they fail to instructor train users because of the inaccessible nature of the knowledge they contain.

Forestry applications. Despite some current work on neural nets for forestry applications (Guan and Gertner 199 lb), few results have yet appeared in the literature. Because neural networks can be used to reformulate empirical mathematical models, it would not be surprising to see neural nets replace their mathematical counterparts for certain applications. Among these applications might be growth and yield models and their various components (e.g., regeneration, mortality, survival, ingrowth). Their two chief advantages over mathematical models are: 1) through their training/learning phase, they implicitly determine the form of the model and estimate its parameters simultaneously, and 2) they can be easily tailored to a specific locale/data by retraining with new, more site-specific data. An obvious criticism of them, however, is that their internal representation of the data (in terms of a model) cannot be explicitly examined, so the model remains hidden.

Because of their ability to perform pattern classification, neural networks can also be useful for classification-type problems. Many forestry problems, in particular many expert system applications in forestry, can be viewed as the categorization of a set of data into a generalized class (Schmoldt 1989). Some examples of classification type problems are: 1) selecting species-site combinations for reforestation, 2) diagnosing insect and disease problems (in fact, all diagnostic problems), 3) selecting silvicultural treatments, and 4) assessing risk. Depending on the needs of system users and on the environment in which the system is deployed, neural networks may be a preferable alternative to expert systems. The former are much easier and faster to develop, but are limited by their black-box opacity.

Future work. Because of the aforementioned advantages, research in neural networks is growing rapidly. However, a number of difficulties still require more research. First, modeling lower-level processes of neurons means that networks must be organized into collective groups to represent higher-level concepts. Second, these higher-level concepts can only be explicitly associated with symbols through the incorporation of language elements. And third, multiple neural nets at varying levels of organization (perhaps a hierarchy) may be necessary to exhibit some of the serial thought processing that is common in high-level reasoning. Hence, some merging of top-down and bottom-up paradigms may produce the range of behavior and thinking with which we are most familiar (Hillman 1990).

Machine Learning

Knowledge of the natural environment exists in the form of corroborated hypotheses, laws, and theories which are all representational models of the natural phenomena we experience. Their status depends on the current state of knowledge and on the confidence with which they are regarded by scientists. When new knowledge is produced, the state, and possibly the form, of these models changes. Thus, knowledge production constitutes learning according to the definition of Michalski (1986): "Learning is constructing or modifying representations of what is being experienced."

We claim that AI techniques can be used to understand and enhance scientific learning or knowledge production. Schank (1987) has stated that, "learning is . . . the quintessential AI issue," and Michalski (1986) claims that "implanting learning capabilities into machines is one of the central goals of artificial intelligence." Developing computational theories of learning and constructing learning systems constitute the subject matter of machine learning. Although current AI systems have very limited learning abilities, it is nevertheless apparent from these systems that rudimentary machine learning capabilities are possible.

Most machine learning research has concentrated on the discovery component of knowledge creation. Michalski (1986) suggests that the approaches to machine learning can be distinguished by the discovery strategies they use to achieve learning. We will discuss the three major strategies, deduction, analogy, and induction, although variations of these and some others also exist. While we are not aware of specific forestry applications of machine learning, the time may be ripe to start investigating how machine learning and other AI methods can help us in conducting our research (Guan and Gertner 199la).

Discovery by deduction. With this strategy, the discoverer draws deductive, truth-preserving inferences from the data and stores them as useful conclusions. One of the first computer programs that was capable of deductively proving theorems, LOGIC-THEORIST (LT), was reported by Newell et al. (1981). The program proved some theorems in the propositional calculus, *Principia Mathematical* (Whitehead and Russell 1913), and has been used as the basis for subsequent systems (0'Rorke 1987). Discovery systems based on the deductive strategy are typically equipped with descriptions of target concepts that are expressed at levels of abstraction too high to be directly usable. The systems use domain knowledge to explain, via a formal proof, why a given fact is an example of a particular concept. Fayyad et al. (1989) have summarized several papers on this topic that were presented at the Fifth International Conference on Machine Learning. One by Braverman and Russell (1989) reports on a system that uses metarules to control the final concepts learned. Another by Rajamoncy and DeJong (1989) considers problems with imperfect domain knowledge which may result in multiple, mutually incompatible explanations.

Discovery by analogy. Discovery by analogy attempts to match descriptions from different domains in order to determine a common substructure which can serve as the basis for the analogical correspondence. Analogical discovery is both deductive and inductive in nature; finding the common substructure involves inductive inference, while performing the analogical mapping is a form of deduction. Examples of systems capable of learning by analogy are described by Winston (1980), Carbonell (1983) and Burstein (1984). Falkenhainer (1987) discusses the use of the analogical strategy to discover and refine qualitative models. The discussion addresses construction of a qualitative theory that goes beyond the discovery of simple empirical laws.

Discovery by induction. Induction is a discovery strategy that draws inferences from the environment. In contrast to deductive systems, correct inputs to an inductive system do not guarantee correct inferences. For a given set of inputs, there is, theoretically, an infinite number of possible inductive inferences. Thus, inductive inference is an underconstrained problem for which one needs additional knowledge to constrain the possibilities and to guide the inference process toward one or a few most plausible hypotheses. Holland et al. (1987) contend that the central problem of discovery via inductive machine learning is specifying processing constraints that will ensure that inferences drawn by the system will be plausible and relevant. Because of the uncertainty of inductive inferences, inductive discovery systems must permit tentative new hypotheses to be included while protecting more certain knowledge from corruption. In realistic situations, logical combinations of tentative hypotheses and background knowledge may produce contradictory results. The system must be able to sort through these contradictory statements, select a coherent subset, and produce an accurate model of the natural environment.

Brief descriptions of some applications using the inductive strategy indicate how they might be adapted to forest science. The symbolic integration system, LEX (Mitchell et al. 1983), develops a general description of a target concept by searching for examples and then including those that are positive, while excluding those that are negative. Researchers at Carnegie-Mellon University have developed AI systems that address several facets of the scientific discovery process (Langley et al 1986). The BACON series of systems focuses on the discovery of empirical laws that summarize numerical data. GLAUBER discovers laws of qualitative structure, such as hypotheses that acids react with alkalis to form salts. STAHL attempts to determine the components of substances involved in chemical reactions and has been used to model reasoning that led to the phlogiston theory. DALTON is concerned with formulating structural models of chemical reactions. While these systems are all interesting in their own right, greater understanding will certainly occur from exploring relations among them and by combining them into a single integrated discovery system.

The future for AI applications to the process of creating knowledge appears bright. Although discovery systems have already been developed, they typically use only one learning strategy and have very limited domains of application. Future discovery systems will be equipped with varying amounts of background knowledge, more or less vague representations of concepts, and theories and laws from other domains. The systems will accept input in the form of examples or observations and will use multiple strategies including analogy, deduction, and induction to generate output in the form of testable hypotheses. These systems will be capable of explaining the results they achieve and the strategy they used to achieve them. In addition, future work will almost certainly focus on separating and abstracting the discovery strategies from the background domain knowledge so that the strategies can be applied to many domains.

Justification. AI applications to enhance the justification component of knowledge creation are not widespread. This component of knowledge creation deals with actually testing hypotheses that have been discovered. Because the truth or falsity of a hypothesis is unknown prior to testing, it is important to consider the kind of evidence necessary for justifying an inference of either kind. In the context of hypothetico-deduction, a deductive machine learning strategy could be applied to the background knowledge with both the hypothesis in question and its negation. For each case, predictions can be generated which can then be compared to empirical evidence. Designing an experiment to simultaneously detect with high probabilities corroborating evidence if the hypothesis is true, and contradicting evidence if it is false, is difficult and complex. Statistical design of experiments has not been ignored by AI researchers, but most of the applications have been in the area of expert systems (Burdick 1987, 1988). While justifying inferences has not received much attention from the AI community, it seems that AI techniques maybe capable of considerable contributions in this area also.

Advisory Systems

Advisory or expert systems are currently among the most visible products in the field of AI. An expert system is a computer program capable of simulating that element of a human specialist's knowledge and reasoning that can be formulated into knowledge chunks so that the computer can approximate the expert's ability to solve problems (Bowerman and Glover 1988). The required knowledge and the inference or reasoning procedures can be thought of as models of the problem-solving expertise of human experts. The ultimate function of advisory systems is to improve the problem-solving skills of non-expert humans (Fig. 3). Thus, an advisory system acts as an organized and accessible repository of the problem-solving knowledge accumulated by human experts. If, as claimed by Feigenbaum et al. (1988), knowledge is the basis of the economic, cultural, and technological power to change circumstances, then advisory systems allow us to automate the problem-solving power of knowledge more effectively.

In the late 1960s and early 1970s, Edward Feigenbaum and Joshua Lederberg developed first expert system, DENDRAL, at Stanford University. It helped determine organic chemical structure using mass spectrometer data. MYCIN, a medical diagnosis and treatment consultant developed by Edward Shortliffe and Bruce Buchanan in the late 1970s, also at Stanford, was another widely publicized expert system. During the 1970s, the conceptual and technical foundations were created for the explosion of expert systems successes in the 1980s. In the 1980s, thousands of advisory systems across many domains were developed (Smart and Knudsen 1986, Walker and Miller 1987). Articles on AI and advisory systems in natural resource management began to appear in significant numbers in 1983 (Davis and Clark 1989). A recent survey article, reporting on 74 projects world-wide (Rauscher and Hacker 1989), showed an increasing number of prototype systems nearing completion, and signaled the emergence of advisory systems as major problem-solving tools in natural resource management. The power and usefulness of advisory systems have now been almost universally recognized (Suzuki 1988).

Many outstanding books on AI and expert systems have been published (Parsaye and Chignell 1988, Walters and Nielsen 1988, Giarratano and Riley 1989, Luger and Stubblefield 1989). Advisory

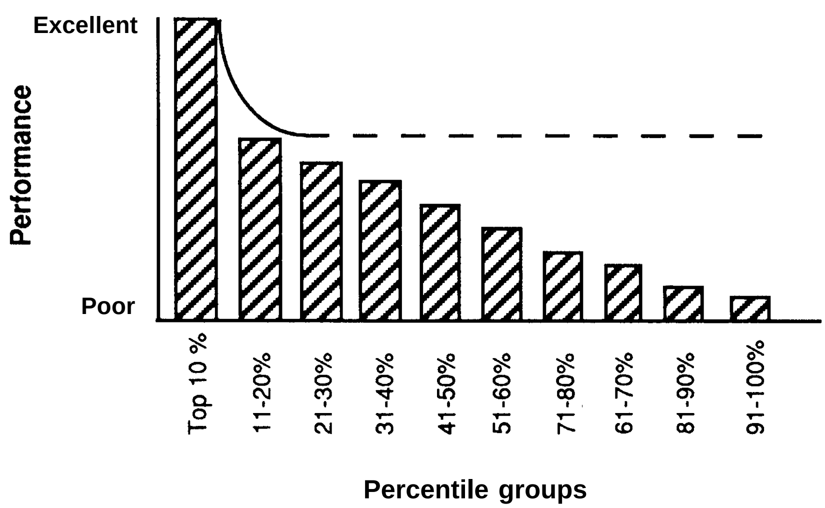

Figure 3. *An expert system can be used to dramatically improve the performance level of non-expert humans. Expert performance is represented by the notched column on the left (top 10%). The expertise level of the advisory system is represented by the dashed line. The expected improvement in performance is the difference between the hatched columns and the dashed line.*

systems applied to natural resource management have been discussed by Starfield (1986) and Schmoldt and Rauscher (in press).

Advisory systems can be used to support the problem-solving component of scientific research by automating *knowledge delivery* and *knowledge organization.* They help identify and organize knowledge relevant to the problem domain and then duplicate and distribute this knowledge as widely as possible. As mentioned earlier, scientists have traditionally published in journals to reach their peers, to archive their work, and to accrue academic recognition. Advisory systems offer an equally powerful mechanism to reach non-academic users, to deliver problem-solving tools to users, and to accrue non-academic recognition from those who ultimately fund research. They provide researchers and research organizations with away to make knowledge a tangible product for the using public. Notice an important change of product here! The research product is not the knowledge embedded in the mind of the researcher, and it is not the result of what can be done with this knowledge, such as a demonstration forest. Rather it is knowledge as a separate entity, a product packaged to help someone solve a problem. This is a new development for forest science. Advisory systems are also cheaper and more convenient because they can tutor individuals at their own locations and at their own pace.

With these advantages, advisory systems are likely to establish their place as forestry problem-solving tools in the next decade. They will be used with increasing frequency by individuals and organizations to distribute knowledge to non-academic users of research results. As the rewards for developing advisory systems become more evident, university and government forestry knowledge centers may be established. These knowledge centers would develop, manage, and expand large, comprehensive advisory systems for a disciplinary specialty, such as boreal forest silviculture. The first products of these centers will likely be small, student-level advisory systems, but with patience, time, and money, these entry-level systems will gradually grow in size and capability until they will form apart of the national knowledge infrastructure.

An example of such evolutionary improvement can be found in the developments of computer chess programs. In 1967, the first computer chess program was developed in the AI laboratory at MIT. While its chess-playing abilities were modest, to say the least, the important thing was that it could play chess, period. After 22 years of evolutionary development, in 1989, the latest computer chess program, developed in the AI laboratory of Carnegie-Mellon University, almost defeated chess grandmaster Anatoly Karpov, the second-ranked player in the world (Steinberg 1990). Similarly, early natural resource expert systems will likely be criticized by human experts, but their quality will almost certainly improve and may even approach the expertise level of human experts.

Knowledge Management

Biological researchers have created an immense body of data and information that is fragmented, unwieldy, and time-consuming to use (Rauscher 1987). If knowledge is defined as organized, evaluated, and synthesized data and information, then the biological sciences are clearly "data rich," but "knowledge poor." Massive amounts of unevaluated and disorganized data and information impede scientific progress and retard cumulative understanding and progress in theory development (Rauscher and Schmoldt 1990). A similar condition has been observed in the social sciences and even in some branches of the physical sciences (Hunter and Schmidt 1990). Knowledge management concepts have evolved out of the AI research effort to advance the development of powerful software tools to cope with this crisis. The cognitive aspects of AI can be viewed as a methodology for "thinking about ways of knowing" (Papert 1988). We are learning to focus explicitly on the importance of knowledge about knowledge (meta-level knowledge). Knowing exactly what it is we do know, determining precisely how we know it, and identifying precisely what we don't know are examples of reasoning about reasoning. A focus on meta-knowledge is the essence of knowledge management systems and is of crucial importance to the future development of the biological sciences.

Scientific progress and organized knowledge. The foundation of science rests on the cumulative growth of knowledge through the application of the scientific method. There are three necessary steps that must function well if scientific progress is to be sustained:

1) primary research studies to test hypotheses;
2) applying meta-analysis tools across primary research studies to establish the knowledge base of facts;
3) applying knowledge management (meta-synthesis) tools to organize the knowledge base of facts into coherent theories of knowledge, thus creating a knowledge base of theory.

For many years, the second step was relatively unimportant, because the number of studies dealing with the same problem was small. This is no longer the case. It is now frequently necessary to resolve differences among a set of studies that all bear on the same relationship. The main focus is to distinguish between variance across studies due to artifacts (such as sampling error or range restrictions) and variance across studies due to real differences (Hunter and Schmidt 1990). The third step has always been recognized as important and has been accomplished by publishing review articles in scientific journals and textbooks. It is now evident that print technology cannot deal effectively with the current volume of scientific knowledge (Rauscher 1987, Rauscher and Schmoldt 1990). The congruence of powerful personal computers, large compact disk electronic storage, and AI-based knowledge management software tools offers a technology powerful enough to deal effectively with the current volume of scientific knowledge.

Efforts have been underway since at least 1945 to devise a knowledge management system capable of managing large masses of information (Parsaye et al. 1989). The basic idea has been to depart from the sequential, linear storage and retrieval of text to a random-access, nonlinear method. The dominant analogue has been the familiar 3 x 5-inch card. The card represents a chunk or node of text, and the trick has been to devise systems that can easily and comfortably structure knowledge by linking these chunks together. The links, which represent the knowledge structure, become as real and as important to the overall system as the chunks, which represent the knowledge content. Such systems are called hypertext systems when the chunks are text, and hypermedia systems when the chunks are graphic images, voice output, or video sequences. The goal is to organize information into systems of knowledge that are intuitive to all users (Larson 1989). Hypertext software systems, enhanced by AI methods, are emerging as powerful knowledge management software tools.

Advisory systems versus knowledge management systems. It is important to recognize the difference between advisory (expert) systems and knowledge management (encyclopedia) systems. According to an essay by Olson et al. (1989), science (pure or basic science) and engineering (applied science) may be placed at either end of a continuum. The essential point to be made is that science is concerned with advancing knowledge, while engineering is concerned with producing and delivering solution-oriented goods and services to decision makers (managers). Engineering may or may not also create new knowledge, but its production is a by-product, not the main objective. Furthermore, new knowledge created as a by-product of the engineering effort tends to be new problem-solving knowledge. Olson et al. (1989) point out that most advisory (expert) systems are clearly solution-oriented tools and, as such, may be labeled more engineering than science. We agree. On the other hand, knowledge management (encyclopedia) systems are concerned with knowledge organization and evaluation in support of further scientific inquiry. Knowledge management software focuses on delivering an encyclopedic knowledge base of theories to scientists and students of science. One typically does not consult an encyclopedia of knowledge to solve a particular problem, but rather to enhance one's understanding. It is, however, recognized that both science and engineering struggle with the "need to order large, complex bodies of information so that rational (scientific and engineering) decisions can be made based on them" (Olson et al. 1989).

A survey of knowledge management systems. Much work has been done to use the power of computers to organize and manage the facts and theories that constitute our social knowledge base. An excellent survey of these efforts, including database systems, bibliographic reference systems, and knowledge management systems, can be found in Parsaye et al. (1989). Rauscher and Host (1990) discuss the intimate connection between hypertext and AI in knowledge management. A good introduction to hypertext methods and techniques can be found in Shneiderman and Kearsley (1989). Rauscher (in prep.) has used hypertext to develop a knowledge management system, the encyclopedia of red pine forest management. An even more ambitious knowledge management project is underway at the Microelectronics and Computer Technology Corporation (MCC). MCC is a consortium of American companies organized to carry out large, high-risk, high-payoff, decade-sized research projects in AI. One of these, called the CYC (encyclopedia) project, hopes to develop a knowledge management system that will contain anon-trivial fraction of the millions of things that we all know as "common sense" and that we assume everyone else knows also (Lemat and Guha 1990). It is this common sense knowledge that people use to know when they don't know and what to do about it.

Knowledge management systems and forest science. In the forest science process, knowledge management systems can be usefully applied to organize, evaluate, and deliver knowledge for the purpose of understanding natural phenomena. Organizing knowledge begins with the explicit identification of facts, hypotheses and laws, and theories. These conceptual components make up the content of the knowledge base. Identifying the same concepts, regardless of the words used to express them, developing a controlled vocabulary that represents them; and linking this controlled vocabulary to occurrences of each concept is a difficult process. First, it is necessary to define each concept, declare its boundaries explicitly, declare its application context and limits, and define what is specifically not included in each concept. Next we need to classify all concepts in ways that are intuitive for other scientists to access these concepts. Several different classifications are usually required in order to access the content of the knowledge base adequately. Finally, we must create structure in the knowledge base by creating links between related concepts. These links turn a collection of information in the form of facts, hypotheses and laws, and theories into a systematically organized body of knowledge. The links may be of many different types of which the following are examples: parent-child, is_a, defines, constrains, contains, special case, example of use. Starting with similar content, experts create highly efficient and useful links, whereas novices more often than not create a mental mess.

Knowledge management systems cannot think for the developer; they only provide powerful methods for implementing the developer's decisions on how to structure knowledge. But once automated into a computerized knowledge base, knowledge management systems allow us to view and evaluate the content and structure of our knowledge explicitly. We can send copies to our peers and debate the additions, deletions, and modifications to both content and structure of the knowledge base. We can identify agreements, as well as disagreements, with an ease never before possible. The knowledge principle from AI states that the secret to intelligent behavior is to have lots and lots of high quality knowledge (Feigenbaum et al. 1988, Lenat and Guha 1990). Computerized database management systems have been accepted as essential aids to the human mind for decades now. No one would dream of trying to manage a large forest inventory on paper or in the minds of humans anymore. Computerized knowledge base management systems are making it equally wasteful to manage forest science knowledge in paper journals and books, or in the minds of our human scientists. The volume is too large and, thanks to advances in AI, the computer can now cheaply store and retrieve knowledge as easily as it can store and retrieve data. What computers cannot do very well is synthesize knowledge creatively, recognize patterns and ascribe meaning to these patterns, know and understand the world around us, and be wise in deciding what we should and should not do to ourselves and our environment.

We are just beginning to explore the power of computerized knowledge management systems. The first step is to develop prototype encyclopedia systems and to discuss their pros and cons. Changes to the links constituting the knowledge structure will become as important as changes to the chunks constituting its content. Forest knowledge management systems, such as the red pine forest management encyclopedia system, are likely to be passed among scientists and re-published each time with improvements. Professional forestry societies may eventually manage and control these changes, with appropriate peer review, and annually issue a revised, accredited version of the entire system. The goal is to create organized, synthesized archives of the cumulative knowledge in forest science that are readily accessible at low costs in time, trouble, and money. Access to such knowledge makes all the difference (Penzias 1989).

Simulation/Modeling

The idea of conserving time and resources by creating artificial representations of reality on the computer has been around for some time (Round 1989). Often this reality consists of numerous interacting components as part of a larger system. When realistic abstractions of reality can be "animated" with a computer model or simulation, valuable insights into the behavior of its real world counterpart can sometimes be gained. However, the combined effect of many components may be difficult to describe by a single mathematical expression. Hence, sets of differential or difference equations have been used in the past to describe the behavior of system components. System behavior of interest may be some current state, some optimal state, or some future state. Certain types of problems, however, may not have any precise mathematical description due to limited and incomplete understanding of the system's underlying mechanisms and, therefore, qualitative relationships must be used. A variety of AI techniques have been applied to these types of problems. These techniques include knowledge-based simulation (Langlotz et al. 1987, Lemmon 1986, Loehle 1987, Meyers and Friedland 1984, Moser 1986, Seliger et al. 1987), common sense reasoning, e.g., qualitative physics (de Kleer and Brown 1984, Forbus 1984, Kuipers 1986), qualitative modeling (Karp and Friedland 1987, Schmoldt 1991a), search (Holland et al. 1986, Nilsson 1971, Pohl 1977), and optimization (Goldberg 1989, Hart et al. 1968).

Knowledge-based simulation. Because these methods cover a number of different topics, we will focus only on knowledge-based simulation and qualitative modeling. By using knowledge-based simulation, one attempts to incorporate heuristic methods into traditional simulation. Here, simulation is used in the sense of a mathematical model of some process that changes over time. Heuristics are added as: 1) a front-end component of the simulator, or 2) an interactive component of the simulator that deals with qualitative aspects of the model. In the former, heuristics infuse some "intelligence" into a simulator by helping the user enter proper input information or by helping the user interpret any output from the simulation. Simulation users may then feel a greater confidence in results, because limitations and assumptions of the models have been adhered to and because results have been explained in a manner more meaningful than tables of numbers.

In the latter case where heuristics appear as an integral component of the simulator, they can be useful for estimating values or selecting among alternatives. Often, heuristics are applied to tasks for which no mathematical solution exists or for which a mathematical solution is expensive or unrealistic. Rauscher et al.'s (1990) work on forest management of red pine simulates establishment, growth, and harvest of red pine stands. Depending on the level at which model heuristics appear, values may be determined heuristically and utilized by a mathematical component, or values may be processed in the reverse fashion from mathematical to heuristic. Most simulations in the past have incorporated qualitative aspects into their models implicitly. However, these judgmental components can now be made explicit and treated in a more rigorous manner.

Qualitative models. Qualitative models have properties that are quite analogous to their quantitative counterparts. Both types of models are completely determined by the values of state variables and the relationships among those variables. Rather than propagating numerical values, however, qualitative models deal with qualitative descriptions of the state variables in a model. In many situations, our understanding of biological processes is incomplete and not well founded in empirical studies. Hence, values like "low,""moderate,""increasing," "steady," etc. are used instead of actual numerical values and relationships. Schmoldt (199lb) used this method for modeling and projecting the combined effects of ozone and drought on mature ponderosa pines. Because variables and relationships are expressed more descriptively, it becomes possible to construct more readily understandable explanations for system behavior. Even in quantitative models, the final output from a model (numerical values) is not used per se but, rather, is interpreted as indicating general relationships or trends (qualitative estimates). Qualitative models incorporate this interpretive step as part of the model specification.

Objects as a modeling paradigm. One of the broadly applicable ideas that has emerged from AI simulation is the use of object-oriented methods. Any process, whether it is a physical process or something like a computer program, can be represented as a collection of interacting objects. These objects interact by communicating with one another via messages that are requests to perform some action. If a particular object has the capability to react to a particular message, it does sousing specialized program code that is part of its own description or that it can legitimately find elsewhere. This paradigm fits very nicely into the task of developing simulations, because all the necessary components of a simulation can be built independently and then plugged into their proper place in the model. This is the basis of several simulation development environments, e.g., Smalltalk (Xerox Learning Research Group 1981), Simula (Birtwistle et al. 1968), and SimKit (Stelzner et al. 1987).

Application to science. Creating a simulator to model some complex process is an integration and synthesis exercise. It forces one to organize the many hypotheses that have been proposed for a particular subject area. Because the hypotheses must fit together as a working unit, a structure, or theory, is imposed on this collection of hypotheses and facts. These hypotheses must talk about the same objects in the same way and must not contain any contradictions. Contradictions are most likely to be noted as the simulation is run. In this way, simulation and modeling help organize knowledge into theories and help evaluate that knowledge as part of a consistent framework. The traditional investigative modes of science, experimentation, and theory previously represented the only means for hypothesizing new knowledge. In the last two decades, a third mode, the computational, has joined the other two, and is rapidly approaching the other two in importance. Results of computational simulation suggest hypotheses beyond the capability of the scientist to generate otherwise.

As a simulation is run on various input data and different assumptions, it may be possible to create new knowledge. Results produced by a model may imply outcomes or emergent properties of the system that were not previously considered. To the extent that these emergent properties are also characteristic of the real world behind the model, new hypotheses may be generated. Results consistent with current understanding might indicate a new hypothesis that could be formulated and tested experimentally. On the other hand, if these outcomes are contradictory with one's expectations, then we might assume that one or more hypotheses of the model are incorrect or, alternatively, the model contains a knowledge gap that causes it to produce erroneous results. In this case, models exist as hypotheses, containing components, interactions, and assumptions that describe current scientific understanding. Simulation of a conceptual model provides justification or negation of a model by producing corroborations or contradictions.

Often such a model of reality contains information that is much different from any publication describing it. A model embodies the theories and hypotheses from which it was built. As such, it allows one to engage in exercises that pose questions to the model and then examine the results for agreement with intuitive expectations. This creates an alternative format to publications for delivering knowledge to other scientists who may then use it to continue the scientific process.

Forestry applications. Modeling and simulation have been used. quite extensively in forestry. Growth and yield models, econometric models, ecological models, operations research models, and forest management models have been simulated on computers. They help us perform experiments that would be prohibitively expensive to actually conduct and they help us envision things that would otherwise be impossible to imagine. Models should become more powerful as heuristic methods extend their flexibility and qualitative models can be used in situations where mathematical models were previously limited by incomplete and inexact data.

Summary

The development of artificial intelligence has been filled with controversy and debate. The search for artificial intelligence is controversial because it touches on the notion of our self-perceived superiority. The ultimate goal of AI practitioners is to create artificial intelligence which resides not inside a computer box, but inside a sensorially capable, autonomously mobile robot. The implications are awesome, while the achievements so far have been much more modest. Whether or not a true artificial intelligence can ever be created is almost beside the point, because the effort to realize this goal has already produced rather significant contributions to computer science, cognitive psychology, philosophy, robotic engineering, and so on.

It should be evident from the previous discussion that AI concepts have a major contribution to make in support of the scientific process. In effect, AI provides methods for implementing epistemological concepts and theories so that they make a tangible and practical contribution to science. The conclusion that has emerged from 40 years of AI research is this: knowledge, not superior reasoning ability, is at the root of intelligence. Knowledge can be treated as a tangible product separate from its storage in the minds of humans.

Knowledge is as real as any physical object. Knowledge is also a theme that runs through most definitions of science as either a body of knowledge or as the creation of knowledge. It is therefore clear that AI and the scientific process are closely allied. Research in AI is calling our attention to the fact that the study of meta-knowledge (knowledge about knowledge) should bean important component of all branches of science. To reiterate: the better we understand how to create, organize, manage, and deliver knowledge, the more efficient we will be as producers, distributors, and consumers of knowledge.

References

Allman, W.F. 1989. Apprentices of wonder. Bantam Books, New York.

Birtwistle, B.M., O.J. Dahl, B. Myhrhaug, and K. Nygaard. 1968. SIMULA BEGIN. Studentliteratur, Lund, Sweden.

Bacon, F. 1620. Novum Organum. London. Republished 1960. New Organon and related writings. Liberal Arts Press, New York.

Barr, A., and E.A. Feigenbaum, editors. 1981. The handbook of artificial intelligence, volume 1. William Kaufmann, Inc., Los Altos, California.

Braverman, M.S., and S.J. Russell. 1989. Boundaries of operationality. Pages 221-234 in: Proceedings, Fifth International Conference on Machine Learning. Morgan Kaufmann Publishers, San Mateo, California.

Bowerman, R,G., and D.E. Glover. 1988. Putting expert systems into practice. Van Nostrand Reinhold Co., New York.

Bunge, M. 1%7. Scientific Research I. Springer-Verlag, New York.

Burdick, R., editor. 1987. Current index to statistics: applications, methods and theory, volume 13. American Statistical Association, Alexandria, Virginia.

Burdick, R., editor. 1988. Current index to statistics: applications, methods and theory, volume 14. American Statistical

Association, Alexandria, Virginia.

Burr, D.J. 1987. Experiments with a connectionist text reader. Pages 717-724 in: Procedngs, IEEE First International Conference on Neural Networks, volume 4, M. Caudill and C. Butler, editors. SOS Printing, San Diego.

Burstein, M. 1984. Concept formation by incremental analogical reasoning and debugging. Pages 351-370 in: Machine Learning An Artificial Intelligence Approach, volume II, R. Michalski, J. Carbonell, and T. Mitchell, editors. Morgan Kaufmann Publishers, Inc., Los Altos, California.

Carbonell, J. 1983. Learning by analogy: formulating and generalizing plans from past experience. Pages 137-162 in: Machine Learning: An Artificial Intelligence Approach, R. Michalski, J. Carbonell, and T. Mitchell, editors. Tioga Publishing Co., Palo Alto, California

Caudill, M. 1989. Using neural nets: representing knowledge. AI Expert 4(12): 34-41.

Davis, J.R., and J.L. Clark 1989. A selective bibliography of expert systems in natural resource management 3(3): 1-18.

de Kleer, J., and J.S. Brown. 1984. A qualitative physics based on confluences. Artificial Intelligence 24:7-83.

Doyle, J. 1979. A truth maintenance system. Artificial Intelligence 12(3): 231-272.

Dretske, F.I. 1986. Knowledge and the flow of information. MIT Press, Cambridge, Massachusetts.

Falkenhainer, B. 1987. Scientific theory formation through analogical inference. Pages 218-229 in: Proceedings, 4th International Workshop on Machine Learning, University of California, Irvine, June 22-25, 1987. Morgan Kaufmann Publishers, Inc., Los Altos, California.

Fayyad, U., J. Laird, and K. Irani. 1989. The fifth international conference on machine learning. AI Magazine 10(2): 79-84.

Feigenbaum, E., P. McCorduck, and H.P. Nii. 1988. The rise of the expert company. Times Books, New York.

Forbus, K.D. 1984. Qualitative process theory. Artificial Intelligence 24:85-168.

Fukumara, T. 1984. Towards the integration of knowledge information processing. New Generation Computing 2: 103-104.

Fukushima, K. 1986. A neural network model for selective attention in visual pattern recognition. Biological Cybernetics 55(l): 5-15.

Giarratano, J., and G. Riley. 1989. Expert systems: principles and programming. PWS-KENT Publishing Company, Boston, Massachusetts.

Giere, R.N. 1984. Understanding scientific reasoning. Holt, Rinehart, and Winston, New York.

Goldberg, D.E. 1989. Genetic algorithms in search, optimization, and machine learning. Addison-Wesley, Reading Massachusetts.

Guan, B.T., and G. Gertner. 1991a. Machine learning and its possible roles in forest science. AI Applications 5(2): 27-36.

Guan, B.T., and G. Gertner. 1991b. Using a parallel distributed processing system to model individual tree mortality. Forest Science (in press).

Halpern, D.F. 1989. Thought and knowledge an introduction to critical thinking. Lawrence Erlbaum Associates, Hillsdale, New Jersey.

Hart, P.E., N.J. Nilsson, and B. Raphael. 1968. A formal basis for the heuristic determination of minimum cost path. IEEE Transactions on SSC 4:100-107.

Hecht-Nielsen, R. 1988. Applications of counterpropagation networks. Neural Networks 1:131-139.

Hillman, D.V. 1990. Integrating neural networks and expert systems. AI Expert 5(6): 54-60.

Hofstadter, D.R. 1979. Gödel, Escher, Bach an eternal golden braid. Vintage Books, New York.

Holland, J., K. Holyoak, R. Nisbett, and P. Thagard. 1987. Induction. MIT Press, Cambridge, Massachusetts.

Hopfield, J.J., and D.W. Tank. 1985. Neural computation of decisions in optimization problems. Biological Cybernetics 52:141-152.

Hume, D. 1739. A treatise of human nature. J. Noon, London. Republished 1978. Oxford University Press, New York.

Hunter, J.E., and F.L. Schmidt. 1990. Methods of meta-analysis: correcting error and bias in research findings. SAGE Publications, Newbury Park, London.

Johnson, G. 1986. Machinery of the mind. Microsoft Press, Redmond, Washington.

Karp, P.D., and P. Friedland. 1987. Coordinating the use of qualitative and quantitative knowledge in declarative device modeling. Knowledge Systems Laboratory Report No. KSL 87-09, Stanford University.

Kleiner, S. 1985. Interrogatives, problems, and scientific inquiry. Synthese 62:365-428.

Kuipers, B. 1986. Qualitative simulation. Artificial Intelligence 29:289-338.

Langley, P., H.A. Simon, G.L. Bradshaw, and J.M. Zytkow. 1987. Scientific discovery. MIT Press, Cambridge, Massachusetts.

Langley, P., J. Zytkow, H. Simon, and G. Bradshaw. 1986. The search for regularity: four aspects of scientific discovery. Pages 425-476 in: Machine Learning: An Artificial Intelligence Approach, volume II, R. Michalski, J. Carbonell, T. Mitchell, editors. Morgan Kaufmann Publishers, Los Altos, California.

Langlotz, C.P., L.M. Fagan, S. Tu, B.I. Sikic, and E.H. Shortliffe. 1987. A therapy planning architecture that combines decision theory and artificial intelligence techniques. Computers and Biomedical Research 20:279-303.

Larson, N. 1989. Hyperlink '89. MaxThink, 44 Rincon Road, Kensington, California.

Lemmon, H. 1986. Comax: An expert system for cotton crop management. Science 233:29-33.

Lenat, D.B., and R.V. Guha. 1990. Building large knowledge-based systems. Addison-Wesley Publishing Company,

Reading, Massachusetts.
Loehle, C. 1987. Applying artificial intelligence techniques to ecological modeling. Ecological Modeling 38:191-212.
Luger, G.F., and W.A. Stubblefield. 1989. Artificial intelligence and the design of expert systems. The Benjamin/ Cummings Publishing Company, Inc., Redwood City, California.
McCarthy, J. 1980. Circumscription-a form of non-monotonic reasoning. Artificial Intelligence 13(1,2): 27-39.
McDermott D.V. 1980. Non-monotonic logic I. Artificial Intelligence 13(1,2): 41-72.
McRoberts, R.E. 1989. A survey of research strategies. Pages 25-31 in: Discovering New Knowledge About Trees and Forests, Selected Papers from a meeting of IUFRO Subject Group S6.09: Philosophy and Methods of Forest Research, August 19-23, 1985, Michigan Technological University, Houghton, Michigan. USDA Forest Service, General Technical Report NC-135.
Meyers, S., and P. Friedland. 1984. Knowledge-based simulation of genetic regulation in bacteriophage lambda. Nucleic Acids Research 12(1): 1-9.
Michalski, R. 1986. Understanding the nature of learning: issues and research directions. Pages 3-24 in: Machine Learning An Artificial Intelligence Approach, volume II, R. Michalski, J. Carbonell, and T. Mitchell, editors. Morgan Kaufmann Publishers, Los Altos, California.
Michie, D., and R. Johnston. 1985. The knowledge machine. William Morrow and Company, Inc., New York.
Minsky, M., and S. Papert. 1969. Perceptions. MIT Press, Cambridge, Massachusetts.
Mitchell, T., P. Utgoff, and R. Banerji. 1983. Learning by experimentation: acquiring and refining problem-solving heuristics. Pages 163-190 in: Machine Learning: An Artificial Intelligence Approach. R. Michalski, J. Carbonell, T. Mitchell, editors. Tioga Publishing Co., Inc., Palo Alto, California.
Moser, J.G. 1986. Integration of artificial intelligence and simulation in a comprehensive decision-support system. Simulation 47(6): 223-229.
Narasimhan, R. 1990. Human intelligence and AI. IEEE Expert 5:77-79.
Newell, A., J.C. Shaw, and H.A. Simon. 1981. Empirical explorations with the logic theory machine: a case study in heuristics. In: Computers and Thought, E.A. Feigenbaum and J. Feldman, editors. Robert E. Krieger Publishing Company, Inc., Malabar, Florida.
Nilsson, NJ. 1971. Problem-solving methods in artificial intelligence. McGraw-Hill, New York.
Olson, R.L., W.P. Kemp, and J.D. Hanson. 1989. AI, expert systems, and science. AI Applications In Natural Resource Management 3:44-48.
O'Rorke, P. 1987. Experimental results of applying explanation-based learning to the logic of Principia Mathematical. Pages 148-159 in: Proceedings, 4th International Workshop on Machine Learning, University of California, Irvine, June 22-25, 1987. Morgan Kaufmann Publishing, Inc., Los Altos, California
Papert, S. 1988. One AI or many? Pages 1-14 in: The Artificial Intelligence Debate: False Starts, Real Foundations, S.R. Graubard, editor. MIT Press, Cambridge, Massachusetts, 1-14.
Parsaye, K., and M. Chignell. 1988. Expert systems for experts. John Wiley and Sons, New York.
Parsaye, K., M. Chignell, S. Khoshafian, and H. Wong. 1989. John Wiley and Sons, New York.
Pearl, J. 1986. On evidential reasoning in a hierarchy of hypotheses. Artificial Intelligence 28(l): 9-15.
Penzias, A. 1989. Ideas and information. Simon and Schuster, Inc., New York.
Pohl, I. 1977. practical and theoretical considerations in heuristic search algorithms. Pages 55-72 in: Machine Intelligence 8, E.W. Elcock and D. Michie, editors. John Wiley and Sons, New York.
Popper, K.R. 1935. Logik der Forschung. In German, republished in English, 1968. The logic of scientific discovery, 2nd edition. Harper and Row, New York.
Rajamoncy, S. A., and G.F. DeJong. 1988. Active explanation reduction: an approach to the multiple explanations problem. Pages 242-255 in: Proceedings, Fifth International Conference on Machine Learning. Morgan Kaufmann Publishers, San Mateo, California.
Rauscher, H.M. 1987. Increasing scientific productivity through better knowledge management. AI Applications in Natural Resource Management 1(2): 21-31.
Rauscher, H.M. In prep. The encyclopedia of red pine forest management. To be submitted to the Northern Journal of Applied Forestry.
Rauscher, H. M., and R. Hacker. 1989. Overview of AI in natural resource management. Pages 202-222 in: Proceedings, AI and Growth Models for Forest Management Decisions, H. Burkhart, H.M. Rauscher, and K. Johann, editors. Vienna, Austria, September 18-22,1989.
Rauscher, H.M., J.W. Benzie, and A.A. Aim. 1990. The red pine forest management advisory system: knowledge model and implementation. AI Applications in Natural Resource Management 4(3): 27-43.
Rauscher, H.M., and G.E. Host. 1990. Hypertext and AI: a complementary combination for knowledge management. AI Applications in Natural Resource Management: 58-61.
Rauscher, H.M., and D.L. Schmoldt. 1990. A role for computers in processing knowledge and in augmenting human thinking. Computer Journal, in press.
Reggia, J.A., B.T. Perricane, D.S. Nau, and Y. Peng. 1985. Answer justification in diagnostic expert systems, Part I. abductive inference and its justification. IEEE Transactions on Biomedical Engineering 32(4): 263-267.
Reichenbach, H. 1938. Experience and prediction. University

of Chicago Press, Chicago, Illinois.
Reiter, R. 1980. A logic for default reasoning. Artificial Intelligence 13(1,2): 81-132.
Round, A. 1989. Knowledge-based simulation. Pages 415-518 in: The Handbook of Artificial Intelligence, volume 4, A. Barr, P.R. Cohen, and E.A. Feigenbaum, editors. Addison-Wesley, Reading, Massachusetts.
Rumelhart, D.E., and J.L. McClelland. 1986. Parallel distributed processing. MIT Press, Cambridge, Massachusetts.
Schank, R. 1987. What is AI, anyway? AI Magazine 8(4): 59-65.
Schmoldt, D.L. 1989. Refining rule bases for classification knowledge-based systems. AI Applications in Natural Resource Management 3(3): 31-41.
Schmoldt, D.L. 1991a An influence model for qualitative simulation of ecological systems. In: Proceedings, 1990 SAF National Convention. Society of American Foresters, 5400 Grosvenor Lane, Bethesda, Maryland. In press.
Schmoldt, D.L. 1991b. Simulation of plant physiological processes using fuzzy variables. In preparation.
Schmoldt, D.L., and G.L. Martin. 1989. Construction and evaluation of an expert system for pest diagnosis of red pine in Wisconsin. Forest Science 35(2): 364-387.
Schmoldt, D.L., and H.M. Rauscher. In press. Building knowledge-based systems for natural resources management. Chapman Hall Publishing Company, New York.
Seliger, G., B. Viehweger, B. Wieneke-Toutouai, and S.R. Kommana. 1987. Knowledge-based simulation of flexible manufacturing systems. Pages 65-58 in: Proceedings, Second European Simulation Multiconference, Vienna, Austria.
Shafer, D. 1976. A mathematical theory of evidence. Princeton University Press, Princeton, New Jersey.
Shneiderman, B., and G. Kearsley. 1989. Hypertext hands-on! an introduction to a new way of organizing and accessing information. Addison-Wesley Publishing Company, Reading, Massachusetts.
Smith, C.G. 1985. Ancestral voices: language and the evolution of human consciousness. Prentice-Hall, Inc., Englewood Cliffs, New Jersey.
Smart, G., and J. Langeland-Knudsen. 1986. The CRI directory of expert systems. Learned Information Ltd., Oxford, England.
Starfield, A.M., and A.L. Bleloch. 1986. Building models for conservation and wildlife management. Macmillan Publishing Company, New York.
Steinberg, D. 1990. Endgame: "Deep Thought" stalks the masters. PC Computing 3:144-149.
Stelzner, M., and T.M. Mitchell. 1985. The SimKit system: knowledge-based simulation and modeling tools in KEE. IntelliCorp, Mountain View, California.
Millings, N. A., M.H. Feinstein, J.L. Garfield, E.L. Rissland, D.A. Rosenbaum, SE. Weisler, and L. Baker Ward 1987. Cognitive science, an introduction. MIT Press, Cambridge, Massachusetts.
Suzuki, N. *1988*. Application specific shells. New Generation Computing 5:317-318.
Tanimoto, S.L. 1987. The elements of artificial intelligence. Computer Science Press, Inc., Rockville Maryland.
Tank, D.W., and J.J. Hopfield. 1986. Simple "neural" optimization networks: an A/D converter, signal decision circuit, and a linear programming circuit. IEEE Transactions on Circuits and Systems 33(5): 533-541.
Walker, T.C., and R.K. Miller. 1986. Expert systems 1987: an assessment of technology and applications. SEAI Technical Publications, Madison, Georgia.
Walters, J. ,and N.R. Nielsen. 1988. Crafting knowledge-based systems: expert systems made realistic. John Wiley and Sons, New York.
Wasserman, P.D. 1989. Neural computing: theory and practice. Van Nostrand Reinhold, New York.
White, H. 1989. Neural network learning and statistics. AI Expert 4(12): 48-52.
Whitehead, A., and B. Russell. 1913. Principia Mathematical. Cambridge University Press, London.
Winston, P. 1980. Learning and reasoning by analogy. Communications of the ACM 23(12): 689-702.
Xerox Learning Research Group. 1981. The Smalltalk-80 System. BYTE Magazine 6(l): 36-47.
Zadeh, L.A. 1965. Fuzzy sets. Information and Control 8:338-353.

Ronald E. McRoberts is a mathematical statistician with the North Central Forest Experiment Station, USDA Forest Service. He received a Ph.D. degree in biostatistics from the University of Minnesota As a statistical consultant, he has collaborated on a wide variety of forestry and forestry-related

projects. He research interests are in nonlinear regression modeling and the logic of scientific research strategies.

Daniel L. Schmoldt is a research technologist with the Southeastern Forest Experiment Station at Virginia Tech. He received degrees in mathematics, computer science, and forest biometry from the University of Wisconsin-Madison. Prior to coming to Virginia Tech, he worked as a research forester for the Fire Management Planning and Economics project and the Atmospheric Deposition project of the Pacific Southwest Research Station. His current responsibilities involve research and development of computer-aided manufacturing systems in the area of primary hardwoods processing. He is also interested in computer vision, pattern recognition, neural networks, and heuristic process control.

H. Michael Rauscher is principal Silviculturist with the North Central Forest Experiment Station. He received his Ph.D. in forest ecology from Virginia Tech and has worked in the area of AI applications in forest management since 1985. He is particularly interested in applications of AI to support the scientific process and expert systems applications in silviculture and ecophysiology. Since 1986, he has been leader of IUFRO Working Party 6.02.01 (Computers in Forestry Research).